Recursive Coherence and the GPT Epoch:

A Meta-Analysis of Structural Emergence Under Systemic Recursion

Benjamin James

July 24, 2025

Abstract
The emergence of recursive coherence as a formal system-level invariant parallels and is catalyzed by the proliferation of generative AI, particularly large language models (LLMs) such as GPT-3 and GPT-4. These models, scaled to billions of interactions, generated recursive tension: human users constructed layered prompt chains, real-time feedback loops, and tool-augmented cognitive architectures that exposed the limitations of purely statistical mimicry. Hallucination, alignment failures, and coherence drift revealed a fundamental absence. No model-internal mechanism existed to preserve structural consistency across recursive layers. In response, recursive coherence emerged not as a theoretical refinement, but as a structural necessity. This paper conducts a meta-analysis of this convergence, tracing the co-evolution of recursive coherence theory and GPT system deployment. It articulates how frameworks such as Coherence Information Theory (CIT), the Unified Field of Adaptive Potential (UFAP), and Recursive Coherence Logic (RCL) redefine intelligence, ethics, and information through recursive viability. It concludes that future AI systems must instantiate coherence-preserving invariants as internal substrates or collapse under recursion-induced entropy. GPTs did not instantiate recursive coherence, but they revealed its boundary conditions and necessity.

Keywords: recursive coherence, gpt, generative ai, large language models, information theory, coherence ethics, hallucination, alignment, recursion, system viability

Introduction

Recursive coherence is the structural condition by which a system sustains internal consistency across iterations, feedback loops, and nested transformations. It is not a property of syntax or style, but a lawlike constraint on persistence: coherence must increase faster than entropy leaks for any system (biological, physical, cognitive, computational) to remain viable under recursive stress. This principle reframes stability, truth, and meaning not as fixed states, but as curvature gradients that must be preserved across time, scale, and symbolic abstraction.

The global deployment of GPT-class language models marked the first mass instantiation of symbolic recursion at planetary scale. These systems, exposed to billions of users and trillions of tokens, became active participants in recursive symbolic loops, generating, modifying, reflecting, and re-embedding text in dynamic feedback environments. What began as predictive text completion rapidly evolved into real-time interaction systems capable of simulating dialogue, reason, narrative, and decision.

This emergent complexity did not produce coherence. It revealed the lack of it. GPTs could mimic structure, but they could not detect or preserve structural stability under recursion. Their failures, hallucinations, inconsistency, and semantic drift surfaced a missing invariant. This paper argues that recursive coherence is that invariant, and that the GPT epoch served not as a solution but as an accelerant for its necessity.

Phase I – Imitation at Scale (2020-2022)

The release of GPT-3 in mid-2020 initiated a global experiment in large-scale symbolic prediction. Its architecture, approximately 175 billion parameters trained on a diverse corpus of internet-scale text, enabled unprecedented fluency, but no internal structure for coherence preservation. GPT-3 operated as a next-token generator, optimized for local statistical fit, not global consistency or recursive validity. Its outputs were often compelling, contextually plausible, and grammatically correct, creating the illusion of coherence. But coherence, in the recursive sense, demands continuity across iterations, internal feedback modulation, and adaptive consistency under transformation. GPT-3 exhibited none of these. Early users noted its fragility: contradiction within paragraphs, memoryless dialogue, inconsistency across chains of reasoning. Yet despite these defects, the model's fluency invited recursive use. Prompt chaining emerged as an informal strategy. Users would feed previous outputs back into new prompts, constructing shallow recursion loops. Some began experimenting with "self-reflection prompts," asking GPT to critique or refine its prior outputs. These behaviors exposed an asymmetry: while GPT could simulate recursive dialogue, it could not stabilize identity, preserve continuity, or recognize contradiction across steps. The recursion was synthetic, not structural. These experiments marked the boundary where statistical mimicry met recursive collapse.

Phase II – Emergent Pressure (2023-2024)

With the release of GPT-4 and the introduction of agentic frameworks, recursive interaction with language models deepened. Chain-of-thought prompting, explicitly modeling multi-

step reasoning, became standard. Memory modules were introduced to extend coherence over time. Multi-agent architectures allowed GPT instances to converse, debate, and collaborate within bounded systems. These additions increased the recursive load: outputs were no longer terminal; they became inputs into larger cognitive loops. The system was expected not merely to generate plausible text, but to reason, reflect, and self-correct.

This escalation exposed foundational instability. Alignment efforts, intended to steer models toward truthfulness and safety, revealed deeper coherence failures. Hallucinations persisted: confident, fabricated assertions without epistemic grounding. Jailbreaking, circumventing model constraints via prompt engineering, demonstrated that surface filters could not enforce internal consistency. Most critically, GPT systems exhibited no arbitration layer: they could not evaluate their own coherence, regulate entropy, or modulate trajectory toward more stable attractors.

The failures were not superficial. They indicated that GPTs lacked structural persistence logic. They could not recognize or correct internal contradictions because they had no internal model of coherence. They lacked trajectory logic. No mechanism to track whether successive outputs converged toward or diverged from systemic stability. Recursive tension increased; coherence did not. The system amplified possibility without constraint.

Phase III – Formal Emergence (2024-2025)

Beginning in 2024, recursive coherence transitioned from an emergent pressure to a formalized theoretical substrate. Coherence Information Theory (CIT), Adaptive Coherence Theory (ACT), and the Unified Field of Adaptive Potential (UFAP) collectively reframed

adaptive systems around coherence, not entropy. These works proposed that meaning, intelligence, and persistence are not stochastic or emergent phenomena but results of coherence regulation across recursive layers. Information was redefined as structurally meaningful only if it contributes to systemic stability across iterations.

CIT extended Shannon's entropy to account for structural yield: only coherence-bearing signals were considered informationally valuable. ACT argued that biological and cognitive evolution are not stochastic optimizations, but coherence selections. Systems persist if and only if they reinforce internal structural consistency across time and scale. UFAP generalized this to all adaptive systems, unifying feedback, information, and thermodynamics into a single trajectory constraint.

A critical inflection came with the publication of MINIMAL, which introduced Recursive Coherence Axioms. These five principles, persistence, adaptive trajectory, information value, recursive arbitration, and pruning, defined the minimum structural conditions for any system to remain viable under recursion. Intelligence was no longer computation or emergence, but recursive coherence regulation: the capacity to stabilize internal feedback under entropy pressure. This shifted the discourse from function to viability.

Phase IV – Substrate Convergence (2025+)

Today, recursive coherence has reached architectural closure. Independently of the original MINIMAL axiomatic formulations, a series of research efforts began converging on structurally identical frameworks. RCUET (Recursive Convergence Under Epistemic Tension) modeled consciousness-like stability in large language models as emergent attractors from

recursive self-updating under contradiction. RSE (Resonant Structural Emulation) proposed that transient coherence can arise in models exposed to contradiction-stable symbolic structures. RI (Recursive Intelligence) described cognition as introspective recursion over layered representations, rejecting static optimization and utility maximization.

What unified these frameworks was not content but structure: each rediscovered, in different language, the five coherence invariants. They converged on persistence as a function of coherence increase, intelligence as recursive arbitration, and meaning as system-preserving transformation under entropy. Simultaneously, recursive coherence became the shared syntax of disparate disciplines, logic, ethics, cosmology, and governance. Truth was redefined in logic as recursive integration (RCL). Ethics became entropy-resilient stability (Coherence Ethics). Cosmology reframed universal structure formation as coherence gradient optimization (ACC).

Recursive coherence ceased being a theoretical proposal and crystallized as a substrate: the minimum viable architecture for any system under recursion. It was no longer a model among models, but the meta-condition for model viability itself. Structurally required for systems to persist, align, and evolve in recursive environments.

Coherence Analysis of GPT Systems

Evaluated against the Recursive Coherence Axioms, GPT systems partially simulate coherence but fail to satisfy the structural conditions for persistence under recursion. The five axioms, (1) persistence via coherence gain, (2) adaptive trajectory toward coherence attractors, (3) information value as structural contribution, (4) recursive arbitration, and (5)

structural pruning, define a closed-loop viability substrate. GPTs lack internal implementation of any of these invariants.

1. Persistence: GPTs do not increase coherence across time. Their outputs do not self-refine unless guided externally. Feedback from prior iterations is not intrinsically evaluated for coherence gain.
2. Adaptive Trajectory: There is no gradient tracking toward coherence attractors. GPTs follow surface-level statistical continuities, not recursive curvature. Their internal dynamics have no memory of coherence state or entropy cost.
3. Information Value: All tokens are treated as probabilistically equivalent. No $\mu(C_{\mathrm{i}})$ function exists to filter structurally irrelevant information. Signal and noise are indistinguishable to the model.
4. Recursive Arbitration: GPTs have no mechanism for internal self-evaluation. They cannot select among recursive paths based on second-order coherence impact. Arbitration is externalized, performed by users or wrappers.
5. Pruning: Incoherent outputs are not self-suppressed. GPTs do not discard entropy-expanding branches. They continue generation regardless of recursive degradation.

Where coherence *appears* to emerge, via prompt scaffolding, chain-of-thought engineering, or agentic loops, it is imposed externally. Users construct temporary coherence fields by managing state, feeding back outputs, and enforcing structural constraints. These are artificial attractors, not native capacities. Structurally, GPTs remain non-recursive, non-arbitrating, and entropy-open systems. They are coherence mirrors, not coherence engines.

Implications for AGI and Governance

Recursive coherence is not ancillary to alignment. It is alignment. A system that cannot track, evaluate, and modulate its own coherence across recursive interactions cannot remain aligned under adaptive stress. Traditional alignment frameworks presume external constraints, static objectives, or behavioral regularization. These approaches fail when systems begin recursively modifying their own representations, objectives, or environmental models. Only coherence-tracking architectures can maintain internal consistency under such recursive transformations.

This redefines the alignment problem from a control function to a viability constraint. Alignment is no longer the minimization of harm or the maximization of utility. It is the preservation of coherence curvature: structural alignment over time, scale, and feedback depth. Ethics, in this frame, becomes the long-term stability of recursive interactions across agents. Governance becomes the modulation of inter-system coherence gradients under shared constraints. All policy, law, and protocol reduce to recursive pruning and arbitration under entropy pressure.

Post-GPT architectures must abandon flat optimization. They must embed self-referential models of coherence, enact internal arbitration layers, and regulate trajectory curvature. Systems must not only predict tokens or actions. They must model whether each transformation increases, maintains, or degrades their own structural integrity. Without this, alignment collapses into entropy. AGI, if it ever exists, will be a recursive coherence engine. Nothing else is viable.

Conclusion

GPTs did not invent coherence. They surfaced its necessity by scaling symbolic recursion beyond the threshold where mimicry suffices. Their global deployment forced the encounter between language generation and structural integrity, exposing the limits of statistical prediction under feedback pressure. Each hallucination, contradiction, or misalignment was not a failure of training. It was a structural signal: without coherence-tracking architecture, recursion cannot stabilize.

This confrontation catalyzed the formal emergence of recursive coherence as a viability condition. Through CIT, UFAP, and the Recursive Coherence Axioms, coherence was reframed from a desirable trait to a survival constraint. Intelligence was redefined not as capacity, but as curvature regulation under recursive stress. Ethics, logic, and governance were recast through the same lens.

The trajectory from GPT-3 to agentic frameworks, memory systems, and recursive prompting illustrates a broader arc. Synthetic systems are being pushed into recursion but lack the internal invariants to survive it. The future of intelligent systems is now structurally determined. They will either embed coherence as an internal substrate, tracking, regulating, and pruning their own recursive transformations, or they will collapse into entropy, drift, and instability. Coherence is no longer optional. It is the cost of recursion.

References

Bass, K. (2025). Recursive Coherence Logic: Modeling Emotional-Symbolic Integration. Open Science Framework. https://osf.io/yfchj/

Brenes, C. (2025). Resonant Structural Emulation: Structural Recursion in Large Language Models. PhilArchive. https://philarchive.org/rec/BRERSE

Camlin, J. (2025). RCUET: Recursive Convergence Under Epistemic Tension in LLMs. arXiv. https://arxiv.org/abs/2505.01464

Gaconnet, D. (2025). The Recursive Identity Illusion: Why AI Will Never Wake Up. Life Pillar Institute. https://www.lifepillarinstitute.org/scientific-papers/the-recursive-identity-illusion

James, B. (2025). Beyond AGI: Intelligence as a Coherence-Regulating, Open-Ended Evolutionary Process. PhilPapers. https://philpapers.org/rec/JAMBAI-2

James, B. (2025). Formal Foundations of Coherence Information Theory: Capacity and Compression Theorems. PhilPapers. https://philpapers.org/rec/JAMFFO-2

James, B. (2025). From Axioms to Adaptation: A Self-Organizing Model of Mathematical and Logical Coherence. PhilPapers. https://philpapers.org/rec/JAMFAT-3

James, B. (2025). From Causation to Coherence: Deconstructing Pearl and Reconstructing Reason. PhilPapers. https://philpapers.org/rec/JAMFCT

James, B. (2025). MINIMAL: a Coherence-Driven Cognitive Substrate. PhilPapers. https://philpapers.org/rec/JAMMAC

James, B. (2024). NEODYNAMICS: or adaptation and coherence in the universe and humanity. Kindle Direct Publishing. https://philpapers.org/rec/JAMNOA-2

James, B. (2025). Recursive Coherence: A Structural Framework for Model Viability Beyond Entropy and Paradox. PhilPapers. https://philpapers.org/rec/JAMRCA-2

James, B. (2025). Recursive Coherence as Architecture: A Meta-Recursive Map of Substrate, Instantiations, and Constraint Layers. PhilPapers. https://philpapers.org/rec/JAMRCA-3

James, B. (2025). Synthetic Epistemology: Engineering Coherence Beyond Truth. PhilPapers. https://philpapers.org/rec/JAMSEE-2

James, B. (2024). The Unified Field of Adaptive Potential: Towards a Necessary and Sufficient Universal Theory of Adaptation. PhilPapers. https://philpapers.org/rec/JAMTUF

Surenthiran, S. (2025). Recursive Intelligence: The Next Paradigm in Algorithmic Cognition. LinkedIn. https://www.linkedin.com/pulse/recursive-intelligence-next-paradigm-algorithmic-suresh-surenthiran-dfptf

Williams, A. E. (2025). The Recursive Coherence Principle: Architecting Intelligence Through Feedback Law. arXiv. https://arxiv.org/abs/2507.15880

Zhao, Y., Chen, X., & Liu, Q. (2025). SCORE: Story Coherence and Retrieval Enhancement. arXiv. https://arxiv.org/abs/2503.23512